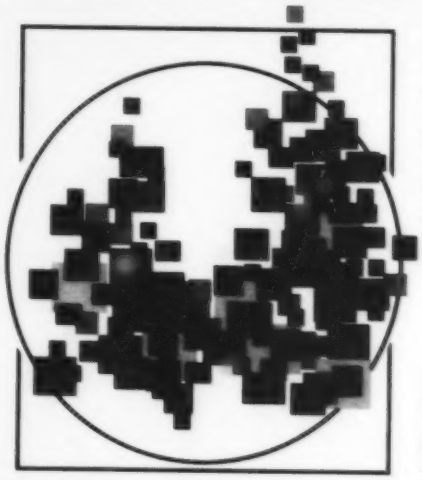

CENTRE FOR
SYSTEMS
SCIENCE

LCCR

Simon Fraser University, Burnaby, B.C., Canada V5A 1S6 / (604) 291-4704

Indeterminate Behavior with Determinate Semantics in Parallel Programs

F. Warren Burton

CSS/LCCR TR 89-03

Indeterminate Behavior with Determinate Semantics in Parallel Programs

F. Warren Burton
School of Computing Science
Simon Fraser University
Burnaby, British Columbia
Canada V5A 1S6

burton@cs.sfu.ca

February 28, 1989

92-05565/

Abstract

A parallel program may be indeterminate so that it can adapt its behavior to the number of processors available, or at least so that low level timing issues are removed from the program.

Indeterminate programs are hard to write, understand, modify or verify. They are impossible to debug, since they may not behave the same from one run to the next.

We propose a new construct, a polymorphic abstract data type called an *improving value*, with operations that have indeterminate behavior but simple determinate semantics. These operations allow the type of indeterminate behavior required by many parallel algorithms.

We define improving values in the context of a functional programming language, but the technique can be used in a procedural programs as well.

1 Introduction

A parallel program may be indeterminate so that it can adapt its behavior to the number of processors available, or at least so that low level timing issues are removed from the program.

For example, in a combinatorial search, many different processes may be searching different subspaces in parallel. These processes may all access and update a global variable that gives information on the best solution found so far. The current value of the variable may be used to determine if a subspace may be pruned from the search. Since the processes are not synchronized, the pruning of a particular subspace may depend on just when a shared variable is read. This causes indeterminate behavior, although the final result may be determinate.

Indeterminate programs are hard to write, understand, modify or verify. They are impossible to debug, since they may not behave the same from one run to the next.

We propose a new construct, a polymorphic abstract data type called an *improving value*, with operations that have indeterminate behavior but simple determinate semantics. These operations allow the type of indeterminate behavior required by many parallel algorithms. Operationally, we may know a lower bound (or upper bound) for an improving value at any given time. If this bound is sufficient, we may act on it. Otherwise, the bound may improve as the computation proceeds.

We define improving values in the context of a functional programming language, but the technique can be used in procedural programs as well. We will user the notation of the Miranda[1] functional programming language [7,8,9].

In section 2 we will briefly review the concept of speculative evaluation. The

[1]Miranda is a trademark of Research Software Ltd.

improving value abstract data type will be introduced in section 3. Several examples of the use of improving values are given in section 4, including a parallel least-cost search algorithm that is only six line long. In these three sections we will limit our attention to two types of computing devices: Those with a single processor and those with an infinite number of processors. Since machines of the second type are not yet on the market, in section 5 we will consider how to make best use of a limited number of processors. The least-cost search algorithm will again be considered. An axiomatic definition of improving values is given in section 6, along with several properties and an outline for a simple correctness proof for the least-cost search algorithm. Section 7 is the conclusion.

2 Speculative Evaluation

Often it is necessary for a good parallel algorithm to perform more work in total than would be performed by a good sequential algorithm for the same problem. This is because the best sequential algorithm may not have a very high potential for parallelism. In other cases, the sequential algorithm may have a high potential for parallelism, but only if some work is done before it is know to be required. We are interested in algorithms of this second kind.

Recall that, by definition, any problem in NP can be solved by a nondeterministic Turing machine in polynomial time. Consider any NP-complete problem, and any nondeterministic polynomial time algorithm to solve the problem where none of the nondeterministic choices lead to a nonterminating computation. A sequential algorithm for this problem can be produced by simulating the nondeterminism by backtracking. In the worst case, we will have to do an exhaustive search using the backtracking algorithm. No other sequential algorithm can do much better in the worst case, assuming $P \neq NP$.

With unbounded parallelism, we can solve any NP-complete problem in polynomial time by considering all of the alternatives of the nondeterministic choices in parallel. Since no sequential algorithm has a worst case polynomial time solution, all NP-complete problems have a high potential for parallelism. On the other hand, the backtracking algorithm could go directly to the solution, solving the problem as quickly as with unbounded parallelism, in the best case. Hence we have an example of a problem where we can gain speed through the use of parallelism, in the average case, but only if we are willing to perform some work that may not be required.

The use of *speculative evaluation* [1,2,3] has been proposed for problems such as this. A speculative computation is a computation that may or may not be required later. For example, while considering one alternative in a backtracking algorithm, speculative computations may be exploring other alternatives.

All computation that is not speculative is called *mandatory*. If a speculative computation is found not to be required, then it may be terminated. If a speculative computation is found to be required, then it must be upgraded to mandatory. Mandatory computation must be favored over speculative computation, at least to the extent that some mandatory computation is always progressing. In the degenerate case of a single processor, no speculative computation will ever be performed. With unbounded parallelism, all speculative computation will be performed, at least until it is found to be not required.

A single function *spec* is sufficient to introduce speculative evaluation. The type of spec is *spec* :: (* -> **) -> (* -> **), and semantically it is the identity function restricted to functions. That is *spec* $f\ x = f\ x$. Operationally, *spec* will initiate the speculative evaluation of its second argument before applying its first argument to its second.

Two useful functions are defined in Fig. 1. The function *spec_or* is defined in

```
spec_or a b = spec (or a) b where or a b = a ∨ b

start [ ] = [ ]
start (x:xs) = spec (spec cons x) (start xs) where cons a b = a:b
```

Figure 1: Two useful functions for initiating speculative computation.

terms of the conditional or operator, ∨. When applied, *spec_or* will evaluate its first argument as a mandatory computation and its second argument as a speculative computation. It the first argument returns *True* then the speculative computation may be terminated by the implementation. If the first argument evaluates to *False* then the speculative evaluation must be upgraded to mandatory if it has not already completed. This function could be useful in a simple backtracking algorithm that returns a boolean result. The function *start* will initiate the speculative evaluation of all of the elements in the list to which it is applied. In Miranda, the colon is an infix cons operator.

Examples of speculative algorithms may be found in [3].

3 Improving Values

Often in combinatorial search algorithms, bounds are maintained to help decide when pruning is possible. For example, in a branch-and-bound search for a least cost solution, a program may keep the value of the best solution found so far. If it is possible to establish that a subspace can not return a better solution, it can be pruned. Alpha-beta search, and similar algorithms, also maintain bounds. In all of these cases, the values of the bounds change monotonically as the computation progresses. We will limit our attention to lower bounds. Upper bounds can be handled in a similar manner.

We proposed the use of the polymorphic abstract data type *improving* for a

```
abstype improving *
with make     :: * -> improving *
     break    :: improving * -> *
     minimum  :: improving * -> improving * -> improving *
     spec_max :: improving * -> improving * -> improving *
```

Figure 2: Signature of type *improving*.

lower bound that may improve (become a tighter bound) over time. The type signature is given in Fig. 2.

We wish to be able to use the result returned by *spec_max* before its second argument has been evaluated if its first argument provides sufficient information, just as we are able to use the result returned by *spec_or* before its second argument has been evaluated provided its first argument is *True*. In general, once the first argument of *spec_max* has been evaluated, we will have a lower bound on the result. If this is not sufficient, we can wait for the bound to improve, which will happen when the value of the second argument is available.

For example, we would like it to be the case that

$$break\ (minimum\ (make\ 5)\ (spec_max\ (make\ 7)\ \bot)) = 5.$$

Here $\bot$ denotes an undefined value or a nonterminating computation. We can't test a value for equality with $\bot$, since the halting problem is undecidable. In practice, we are not really concerned with handling nonterminating computations or undefined values. However, if any subexpression of an expression can be replaced by $\bot$ without changing the value of the expression, then the subexpression need not be evaluated.

In practice, it would be better to parameterize the *improving* abstract data type with an order relation. This would permit *improving* values to be used for

upper bounds as well as lower bounds. This would also support applications that require order relations other than those build into the language. This change is straight forward, but complicates the presentation slightly, so will be omitted. In this paper, for simplicity, we will rig all of our examples so that the "<" operator, which is provided for all types in Miranda, gives the desired ordering.

An improving value can be represented by a strictly increasing list of approximations. The final element of the list will be the true value that previous list elements approximates. Infinite lists and partial lists will be considered in section 6. We will assume that for any two values a and b, the value of $a < b$ is defined.

With this representation, the *improving* abstract data type can be implemented as shown in Fig. 3. The function *make* produces a singleton list, and *break* returns the last element of a list. The *spec_max* function starts the speculative evaluation of its second argument and then appends its arguments, removing any values in the second list that are less than or equal to the final element of the first list. This insures that the list of approximations remain strictly increasing. On the other hand, *minimum* merges two lists, removing duplicates in order to maintain strict monotonicity. In addition, the merge ends as soon as either list ends. For example,

minimum (*make* 4) (*spec_max* (*make* 3) (*make* 5))

will cause the lists [4] and [3, 5] to be merged to produce [3, 4], with the 5 thrown away. As a second example,

minimum (*make* 4) (*spec_max* (*make* 5) (*make* 3))

will result in the two singleton lists [4] and [5] being merged to produce the singleton list [4]. The 3 will be discarded by *spec_max* since it would be out of order, and the 5 will be discarded by *minimum* as before.

```
improving * == [*]

make a = [a]

brak x = last x

spec_max xs ys = spec (monotonic_append xs) ys

minimum xs ys = short_merge xs ys

monotonic_append [x] ys = x : dropwhile (<= x) ys
monotonic_append (x:xs) ys = x : monotonic_append xs ys

short_merge xs [ ] = [ ]
short_merge [ ] ys = [ ]
short_merge (x:xs) (y:ys)
   = x:short_merge xs ys, x = y
   = x:short_merge xs (y:ys), x < y
   = y:short_merge (x:xs) ys, x > y
```

Figure 3: A simple implementation of *improving*.

We will sometimes refer to the elements of a list representing an *improving* value as a *progress report*. Each progress report gives some new information about the final value of an *improving* value.

It is easy to see that every element of a list representing an *improving* value, except for the final element of the list, must at some time have occurred in a left argument of an application of *spec_max*. This property is trivially true of the singleton lists produced by *make*, and is preserved by *minimum* and *spec_max*. Often it is useful to use the expression *spec_max a b* in situations where *b* is the desired result, but *a* is known to be a lower bound on the result. In this case, *a* is a progress report.

If the abstract type *improving* is implemented as a language primitive, then a more efficient implementation than the one described above is possible. Each *improving* value is represented by a pair consisting of the best approximation found so far and a flag indicating whether this is a final value. Each *improving* value has an associated process that will update the approximation as required and supply the latest value to other processes upon request. This save storing more than one value and saves other processes the cost of examining out of date values before coming to the most recent value. In some cases this can significantly reduce the amount of communication between processes.

4 Examples

In this section we will consider parallel versions of three search algorithms.

4.1 Least-Cost Search

With many combinatorial problems it is necessary to search a solution space for the best solution to the problem. We will assume that "best" means "smallest"

and call the measure of a solution the *cost* of the solution. Horowitz and Sahni present a least-cost search algorithm for this problem [5]. Assuming the costs of leaves and lower bounds for other nodes are all distinct, this algorithms expands exactly the same nodes as the A^* algorithm in the case of a single processor, which is optimal [6].

We will assume that the solution space is organized as a tree with all possible solutions at leaves. For any leaf node, *leaf*, *cost leaf* is the cost of the solution. For any node, *is_leaf node* returns a *bool* indicating whether the node is a leaf. If the node is not a leaf, then *children node* returns a list of the children of *node*. In addition, if *node* is not a leaf, *lower_bound* gives a lower bound on the least cost solution to be found in the subtree rooted at *node*. Finally, let us assume that *nil_node* is a special node such that *nil_node* < *node* for any node, *node* that may be encountered during a search. (Otherwise nodes are ordered arbitrarily.)

Let us first consider a simple exhaustive search algorithm to solve this problem. Such an algorithm is given in Fig. 4. The function returns an ordered pair, consisting of the cost of the best solution together with that solution. Note that if *cost1* < *cost2* then (*cost1*, *node1*) < (*cost2*, *node2*) in Miranda. We should note that "." is an infix function composition operation and *fold* is a function that will apply a binary function to elements of a list, reducing the list to a single value (e.g. *fold* (+) *xs* will compute the sum of the elements of the list *xs*). We will use *mini* and *maxi* for binary minimum and maximum functions, respectively. The function *map* will apply a function individually to elements of a list, producing a new list.

Fig. 5 gives an equivalent (but slightly less efficient) version of this search. Since we know that *lower_bound root* is a lower bound on the cost of the node returned by searching the subtree rooted at *root*, and in case it is a tight lower bound *nil_node* is less than any other node, it follows that

```
search root
    = (cost root, root), is_leaf root
    = (fold mini.map search.children)root, otherwise
```

Figure 4: An exhaustive search algorithm.

```
search root
        = (cost root, root), is_leaf root
        = max (lower_bound root, nil_node)
          ((fold mini.map search.children) root), otherwise
```

Figure 5: A modified exhaustive search algorithm.

$$(lower_bound\ root,\ nil_node) < (fold\ mini.map\ search.children)\ root$$

so the two algorithms must return the same result.

A minor further modification, to introduce *improving* values, yields the least-cost search algorithm in Fig. 6. Notice how the search of each subtree starts by providing a progress report in the form of a lower bound on the best solution to be found in the subtree. If a better solution has been found anywhere in the tree, this progress report is sufficient for the tree to be pruned. In the case of a single processor, where no speculative evaluation is performed until it becomes necessary, only those nodes with a lower bound less than or equal to the cost of the optimal solution can ever be expanded. With an unbounded number of processors, all paths are searched in parallel, at least until a least cost solution is found.

```
search = break.search'
         where
         search' root
            = make (cost root, root), is_leaf root
            = spec_max (make (lower_bound root, nil_node))
                ((fold minimum.map search'.children) root), otherwise
```

Figure 6: A least-cost search algorithm.

4.2 Breadth-First Search

Fig. 7 shows a parallel breadth-first search algorithm. The algorithm takes a node as a parameter and returns a pair consisting of the depth of a least deep solution and the solution itself. We assume that we are given two functions: *is_solution*, which determines whether a node is a solution to the problem of interest, and *children*, which will generate a list of the children of a node. We do not require that a solution be a leaf node. As with the least-cost search algorithm, we assume the existence of a node *nil_node*. We also assume a value *infinity* which is greater than the depth of any reasonable search. This simplifies the algorithm and allows us to use the built in order relation "<". Notice that we have used the function *start* defined in Fig. 1 to initiate the parallel searching of all subtrees when the subtree root has been found not to be a solution.

4.3 Alpha-Beta Search

As a final example we will consider a parallel alpha-beta search algorithm for searching a game tree, based on the sequential alpha-beta search algorithm given in [5]. For simplicity, the algorithm returns the value of the best move,

```
breadth_first = break.search 0
search depth root
   = make (depth, root), is_solution root
   = make (infinity, nil_node), kids = [ ]
   = spec_max progress_report solution, otherwise
     where
     kids = children root
     progress_report = make (depth, nil_node)
     solution = fold minimum (start [search (depth + 1) k | k <- kids])
```

Figure 7: A parallel breadth-first search algorithm.

not the move itself. Again *start* is used to initiate the searching of subtrees. The algorithm is given in Fig. 8.

We assume the existence of three functions. For leaf nodes *eval* computes the value of the node to the player whose turn it is, and for other nodes *children* computes the children of the node (of which we assume that there is at least one.) Finally, *is_leaf* determines whether a node is a leaf.

The function *scan* is a commonly used function, similar to *fold*, but returning a list of "partial sums" rather than just the final "sum". It is defined by

```
scan f a xs
   = a : scan' f a xs
     where
     scan' f a [] = []
     scan' f a (x:xs) = scan f (f a x) xs
```

The function *zip2* maps two lists into a list of corresponding pairs, ending as soon as either list ends.

```
alpha_beta root alpha beta
    = eval root, is_leaf root
    = break (minimum (make beta) best), otherwise
      where
      best = (fold spec_max.map make) alphas
      alphas
          = spec (scan maxi alpha)
                 (start [ - (alpha_beta child ( - beta) ( - new_alpha)) |
                          (child, new_alpha) <- zip2 (children root) alphas])
```

Figure 8: An alpha-beta search algorithm.

The expression (*alpha_beta node alpha beta*) searches for the value of the best move for the player whose turn it is, subject to the constraint that only moves with value between *alpha* and *beta* are considered. We know that by making a different move, we can get to a position with value at least *alpha*, and also know that our opponent can keep us from getting to a position of value greater than *beta* by making a different move earlier. The initial call is of the form

$$alpha_beta\ root\ (-\infty)\ \infty$$

The list *alphas* is the list of alpha values that result after the search of each child.

Parallel algorithms for minimax searching is an active area of research. This simple parallel alpha-beta algorithm is probably not the best solution to the problem. Our notation makes it easier to understand and verify algorithms, but fundamental problems of finding the best algorithm for a given problem remain.

5 Priorities

If all computation is mandatory, then scheduling is not a difficult problem, assuming a shared memory of sufficient size, so we do not need to worry about communication between processors or running out of memory. There are simple scheduling algorithms [4] that are, in the worst case, within a factor of two of being optimal, where we measure the quality of a scheduling algorithm by the time required to finish all computation.

This is not true with speculative computation. If we have n processors and a potential for parallelism that is much higher than n, then it is possible for some problems to get a speedup approaching n if all processors do mandatory work or speculative work that will later become mandatory almost all of the time. On the other hand, if one processor does mandatory work and all other processers spend almost all of their time on speculative computation that will prove to be unneeded, then the speedup may not be much greater than one. Clearly, given a limited number of processors, we prefer to do that speculative work that is most likely to be required later. In general, if is not possible for the implementation to determine which speculative computations are the most worthwhile.

The solution to this problem is to let the programmer specify priorities. We can introduce a new function, *priority* of type num -> * -> *. The semantics of *priority* are

priority $n\ x =$ **if** $n = \perp$ **then** $\perp$ **else** x

If priority is called within a speculative computation, it initiates a new speculative computation with priority n, where n is any number, to compute x. The higher the value of n, the higher the priority of the speculative computation. All speculative computations started with *spec* have higher priority than any started

by *priority*. The priority of a speculative computation can not be changed, except that the speculative computation may become mandatory. In particular, when one speculative computation finds it needs the result of another, possibly lower priority, speculative computation in order to proceed, it must wait for the second speculative computation to finish (unless, of course, the waiting computation becomes upgraded to mandatory.) Of course, if *priority* is invoked by a mandatory computation, then the resulting computation will immediately become mandatory, since it would not have been invoked if its result were not needed.

As an example, we could modify the least-cost search algorithm in Fig. 6 by changing the subexpression

(*fold minimum.map search'.children*) *root*

to

(*priority* (1/*lower_bound root*).*fold minimum.map search'.children*) *root*

so that nodes that are more promising (have a smaller lower bound on the least cost solution) are expanded with higher priority. (Of course, we would want to factor out the computation of *lower_bound root* and compute it only once.)

6 Formal Properties of Improving Values

We will assume that the values used in improving values are finite in size and come from a flat domain. That is, values are either completely defined or are $\perp$. An example of a nonflat domain would be the domain of lazy lists. For example, the lazy list $1{:}\perp$ has 1 as its first element, but any attempt to evaluate the tail of the list will result in a nonterminating computation or an undefined result.

If we allow values from a nonflat domain to be improving values, we have some awkward special cases. For example, with lists ordered in the obvious way,

$$break\ (minimum\ (make\ (1{:}\bot))\ (spec_max\ (make\ (1{:}\bot))\ (make\ (2{:}\bot))) = \bot,$$

since the comparison of $(1{:}\bot)$ with $(1{:}\bot)$ will not terminate. On the other hand,

$$mini\ (1{:}\bot)\ (maxi\ (1{:}\bot)\ (2{:}\bot)) = 2{:}\bot.$$

In this case, *maxi* is applied first and the nonterminating comparison does not arise. Similar problems can arise with infinite lists. We would like the improving value operations to be at least as well defined as the corresponding operations on ordinary values.

Sometimes it will be useful to have improving values where values come from a nonflat domain. For example, we might want to use improving lists in a context where we know that all lists will be finite and well defined. In section 4, we used improving ordered pairs. The domain of ordered pairs is not flat, since it includes elements such as $(1, \bot)$. However, we only used fully defined ordered pairs. In cases such as these, we must restrict ourselves to a flat subdomain. If type *improving* is implemented as a primitive in a language, the restriction to finite, fully defined values can be enforced by requiring *make* to fully evaluate its argument. The results in the remainder of this section depend on values being finite and either fully defined or completely undefined before *make* is applied.

With the implementation of type *improving* given in Fig. 3, a number of different list may all represent the same abstract object. Since *break*, *spec_max* and *minimum* all examine list elements sequentially, lists of the form xs ++ $[\bot]$ ++ ys, where "++" is an infix append operator, can not be distinguished from lists of the form xs ++ $[\bot]$. Also, since lists are strictly increasing sequences of improving approximations, if a and b are any two well defined values with $a < b$, then it is not possible to distinguish $a{:}b{:}xs$ from $b{:}xs$ provide that $b{:}xs$ is a valid representation for an *improving* value.

With these observations, we can divide representations into equivalence classes. We will let $a!$ represent the class of all finite, fully defined lists with final

value a. If xs is a member of the equivalence class $a!$, then we will represent the equivalence class that includes xs ++ [⊥] by $a+?$. We will read $a!$ as "exactly a", and $a+?$ as "at least a". Both [⊥] and ⊥ belong to the same equivalence class which we will represent by ⊥ and call "bottom". Finally, we have two cases to consider with infinite lists. (We should note that if numbers are represented to a finite precision, infinite lists of strictly increasing values can't occur.) The class of infinite lists with no upper bound on their elements is represented by ∞, and an infinite list where the elements have a least upper bounded a is a member of the equivalence class represented by a-ϵ, which is read "almost a".

Axioms for *improving* values are given in Fig. 9. In this figure, a and b may represent any values and x may represent any *improving* value.

With these axioms, we can prove a number of interesting properties. These include:

$$make\ (mini\ a\ b) = minimum\ (make\ a)\ (make\ b) \tag{1}$$

$$make\ (maxi\ a\ b) \sqsubset spec_max\ (make\ a)\ (make\ b) \tag{2}$$

$$mini\ (break\ a)\ (break\ b) \sqsubset break\ (minimum\ a\ b) \tag{3}$$

$$maxi\ (break\ a)\ (break\ b) = break\ (spec_max\ a\ b) \tag{4}$$

$$break.make = id \tag{5}$$

$$make.break \sqsubset id \tag{6}$$

We can use these properties to prove the correctness of efficient combinatorial search algorithms using *improving values*. To show that an implementation meets its specification, we show that *specification* ⊏ *implementation*. That is, where the specification is defined, the implementation must agree with it, but the implementation may be stronger. For example, we may take an exhaustive

$make\ \bot = \bot$

$make\ a = a!$

$break\ \bot = \bot$

$break\ a! = a$

$break\ a+? = \bot$

$break\ a\text{-}\epsilon = \bot$

$break\ \infty = \bot$

$minimum\ \bot\ x = \bot$

$minimum\ x\ \bot = \bot$

$minimum\ \infty\ x = x$

$minimum\ x\ \infty = x$

$minimum\ a!\ b! = \textbf{if}\ a < b\ \textbf{then}\ a!\ \textbf{else}\ b!$

$minimum\ a!\ b+? = \textbf{if}\ a < b\ \textbf{then}\ a!\ \textbf{else}\ b+?$

$minimum\ a!\ b\text{-}\epsilon = \textbf{if}\ a < b\ \textbf{then}\ a!\ \textbf{else}\ b\text{-}\epsilon$

$minimum\ a+?\ b! = \textbf{if}\ a < b\ \textbf{then}\ a+?\ \textbf{else}\ b!$

$minimum\ a+?\ b+? = \textbf{if}\ a < b\ \textbf{then}\ a+?\ \textbf{else}\ b+?$

$minimum\ a+?\ b\text{-}\epsilon = \textbf{if}\ a < b\ \textbf{then}\ a+?\ \textbf{else}\ b\text{-}\epsilon$

$minimum\ a\text{-}\epsilon\ b! = \textbf{if}\ a \leq b\ \textbf{then}\ a\text{-}\epsilon\ \textbf{else}\ b!$

$minimum\ a\text{-}\epsilon\ b+? = \textbf{if}\ a \leq b\ \textbf{then}\ a\text{-}\epsilon\ \textbf{else}\ b+?$

$minimum\ a\text{-}\epsilon\ b\text{-}\epsilon = \textbf{if}\ a < b\ \textbf{then}\ a\text{-}\epsilon\ \textbf{else}\ b\text{-}\epsilon$

$spec_max\ \bot\ x = \bot$

$spec_max\ \infty\ x = \infty$

$spec_max\ a+?\ x = a+?$

$spec_max\ a\text{-}\epsilon\ x = a\text{-}\epsilon$

$spec_max\ a!\ \bot = a+?$

$spec_max\ a!\ \infty = \infty$

$spec_max\ a!\ b! = \textbf{if}\ a < b\ \textbf{then}\ b!\ \textbf{else}\ a!$

$spec_max\ a!\ b+? = \textbf{if}\ a < b\ \textbf{then}\ b+?\ \textbf{else}\ a+?$

$spec_max\ a!\ b\text{-}\epsilon = \textbf{if}\ a < b\ \textbf{then}\ b\text{-}\epsilon\ \textbf{else}\ a+?$

Figure 9: Axioms for type *improving*.

search algorithm as a specification for a more efficient search algorithm that avoid searching unnecessary subspaces. In these cases, the implementation must exceed the specification, because nonterminating computations in the pruned portion of the search space will be unencountered by the implementation, but would cause the specification to fail to terminate.

If we take the algorithm in Fig. 4 as a specification for the more efficient least-cost search algorithm of Fig. 6, then we can easily prove the least-cost search algorithm correct. First we recall that the algorithm in Fig. 5 is equivalent to the one in Fig. 4. Using properties 6, 1, 2 and 5 above, it is easy to show that the least-cost search algorithm meets its specification.

7 Conclusion

We have seen that the polymorphic abstract data type *improving* allows us to express various combinatorial algorithms in a manner that is simplier than most current expressions of such functions, yet at the same time introduces parallelism into the problem. Furthermore, the type *improving* has an axiomatic specification from which we can derive several important properties, which in turn can be use to prove the correctness of programs using the type.

As an example, we presented a six line function for a parallel least-cost search algorithm that is optimal on a single processor and can make good use of any number of processors available. A correctness proof of the function was outlined and can be easily completed by the reader.

References

[1] F. Warren Burton. Controlling speculative computation in a parallel functional programming language. In *Proceedings of The Fifth International*

Conference on Distributed Computing Systems, pages 453–458, Denver, Colorado, May 1985.

[2] F. Warren Burton. Functional programming for concurrent and distributed computing. *Comp. J.*, 30(5):437–450, Oct. 1987.

[3] F. Warren Burton. Speculative computation, parallelism, and functional programming. *IEEE Trans. Comput.*, C-34(12):1190–1193, Dec. 1985.

[4] Derek L. Eager, Edward D. Lazowska, and John Zahorjan. Speedup versus efficiency in parallel systems. *IEEE Trans. Comput.*, to appear.

[5] Ellis Horowitz and Sartaj Sahni. ***Fundamentals of Computer Algorithms.*** Computer Science Press, 1978.

[6] Nils J. Nilsson. ***Problem-Solving Methods in Artificial Intelligence.*** McGraw-Hill, 1971.

[7] David A. Turner. Functional programs as executable specifications. In C. A. R. Hoare J. Shepherdson, editor, ***Mathematical logic and programming languages***, pages 29–54, Prentice-Hall, 1985.

[8] David A. Turner. Miranda: a non-strict functional language with polymorphic types. In Jean-Pierre Jouannaud, editor, ***Functional Programming Languages and Computer Architecture, Lecture Notes in Computer Science, 201***, pages 1–16, Springer-Verlag, 1985.

[9] David A. Turner. An overview of Miranda. *SIGPLAN Notices*, 21(12):158–166, Dec. 1986.

Centre for Systems Science

• Simon Fraser University established the **Centre for Systems Science** as a provincial centre for advanced study. The objective of CSS is twofold: coordinate and combine advanced systems research at SFU, and encourage application of research results within industry in order to further the technological development of British Columbia.

•CSS has assembled three research groups to streamline its efforts. The orientation of these groups is as follows: intelligent systems, microelectronics, and computer and communications systems. Because of shared research interests, this latter group merges with LCCR (**Laboratory for Computer and Communications Research**) for the publication of technical reports.

• CSS's three groups involve over 80 skilled faculty researchers principally from SFU's Schools of Engineering Science and Computing Science, but also draw substantially on Mathematics, Physics, Geography, Philosophy, Linguistics, Psychology, Communication, and Kinesiology.